arXiv:1609.06522v1 [cs.DM] 21 Sep 2016

Computing Vertex-Disjoint Paths using MAOs

Johanna E. Preißer*
Institute of Mathematics
TU Ilmenau

Jens M. Schmidt*
Institute of Mathematics
TU Ilmenau

Abstract

Let G be a graph with minimum degree δ. It is well-known that maximal adjacency orderings (MAOs) compute a vertex set S such that every pair of S is connected by at least δ internally vertex-disjoint paths in G.

We present an algorithm that, given any pair of S, computes these δ paths in linear time $O(n+m)$. This improves the previously best solutions for these special vertex pairs, which were flow-based. Our algorithm simplifies a proof about pendant pairs of Mader and makes a purely existential proof of Nagamochi algorithmic.

1 Introduction

Vertex-connectivity is a fundamental parameter of graphs that, according to Menger [11], can be characterized by the existence of internally disjoint paths between vertex pairs. Thus, much work has been devoted to the following question: Given a number k, a simple graph $G = (V, E)$, and two vertices of G, compute k internally disjoint paths between these vertices if such paths exist. Despite all further efforts, the traditional flow-based approach by Even and Tarjan [3] and Karzanov [6] gives still the best bound $O(\min\{k, \sqrt{n}\}m)$ for this task, where $n := |V|$ and $m := |E|$.

Our research is driven by the question whether k internally disjoint paths can be computed faster. We have no general answer, but show that for specific pairs of vertices, this can actually be done using *maximal adjacency orderings (MAOs)*. MAOs order the vertices of a graph, are also known under the name *maximum cardinality search*, and can be computed in time $O(n+m)$ [17]; we give a definition of MAOs later. On a side note, we exhibit apparently forgotten similarities of MAOs to a very early proof by Mader about pendant pairs in 1971 [8, 7]. In fact, MAOs (and maximum cardinality search) can be seen as a simpler version of Mader's result.

*This research was supported by the DFG grant SCHM 3186/1-1.

Consider a simple graph G with minimum degree δ and a MAO $<$ of G. Let a subset S of vertices be *k-connected* if G contains k internally vertex-disjoint paths between every two vertices of S. Henzinger proved for every $1 \leq k \leq \delta$ that the last $\delta-k+2$ vertices of $<$ are k-connected [5]. Nagamochi [12] improved this result to the more general vertex sets of trees derived from the *forest decomposition* of the MAO; one of these trees actually contains the last $\delta - k + 2$ vertices given above. Nagamochi's proof uses the machinery of mixed connectivity, which generalizes both edge- and vertex-connectivity.

If we desire δ-connected vertex sets, this may result only in a set of size two in the worst case. However, in practice, many more of these sets may occur and each of them may have a much larger size. Even in the worst-case, the approach allows us to increase the size $\delta - k + 2$ of these sets at the expense of linearly decreasing their connectivity to any prescribed $k < \delta$.

We consider the problem of computing k internally vertex-disjoint paths between any vertex pair $\{s, t\}$ that is contained in one of the above k-connected sets. By the above results, we already know that these paths exist. However, the proofs of Henzinger and Nagamochi are non-constructive and therefore do not give any faster algorithm for finding these paths than the traditional flow-based one.

We give an algorithm that computes these k paths in linear time $O(n + m)$. The algorithm proceeds by a right-to-left sweep in the MAO through the vertices of the corresponding set. We also show how the computation can be extended to find the k internally vertex-disjoint paths between a vertex and a vertex set, and between two vertex sets, whose existence was shown by Menger [11].

Certifying Algorithms. Being able to compute k internally vertex-disjoint paths has a benefit that purely existential proofs or algorithms that only argue about vertex separators do not have: It certifies the connectivity between the two vertices. This has been used for *certifying algorithms*, as proposed in [10].

E.g., the famous *min-cut* algorithm of Nagamochi and Ibaraki [13, 4], as simplified by Stoer and Wagner [16] (and initiated by Mader) uses the k-connected vertex sets above by just setting $k := \delta$ and contracting the last two δ-connected (and thus δ-edge-connected) vertices of a MAO iteratively. For unweighted multigraphs, this can be easily made certifying: The certificate just stores the k edge-disjoint paths between the last two vertices $\{s, t\}$ of the MAO in every step; the global k-edge-connectivity then follows by transitivity. In fact, the desired k edge-disjoint paths are given directly by considering the (edge-disjoint) trees $T_1, \ldots, T_k$ of the MAO that contain t and taking the unique s-t-path in each of them.

Using more involved methods, Arikati and Mehlhorn [1] showed that the Nagamochi-Ibaraki algorithm can be made certifying even for weighted graphs,

again without increasing the asymptotic running time and space. For the problem of testing a graph on being k-(vertex-)connected, finding a good certificate seems to be an open graph theoretic problem, even when k is fixed:

Open Problem. For $k \in \mathbb{N}$ and a simple graph G, find a small and easy-to-verify certificate that a graph is k-connected.

So far, linear-time certifying algorithms are only known for $k \leq 3$ [18, 15]. For $k \geq 4$, the best known certifying algorithm still seems to be the traditional flow-based one [3], which achieves a running time of $O((k+\sqrt{n})k\sqrt{n}m)$. Should at some day the best algorithm for computing vertex-connectivity be based on MAOs, our data structure provides the means to turn this into an efficient certificate.

2 Maximal Adjacency Orderings

Throughout this paper, our input graph $G = (V, E)$ is simple, unweighted and of minimum degree δ. We assume standard graph theoretic notation as in [2]. A *maximal adjacency ordering* (MA Ordering; MAO) $<$ of G is a total order $1, \ldots, n$ on V such that, for every two vertices $v < w$, v has at least as many neighbors in $\{1, \ldots, v-1\}$ as w has. For ease of notation, we always identify the vertices of G with their position in $<$.

Every MAO $<$ partitions E into forests $F_1, \ldots, F_m$ (some of which may be empty) as follows: If $v > 1$ is a vertex of G and $w_1 < \cdots < w_l$ are the neighbors of v in $\{1, \ldots, v-1\}$, the edge $\{w_i, v\}$ belongs to F_i for all $i \in \{1, \ldots, l\}$. For every i, the graph (V, F_i) is a forest (we refer to [14, Section 2.2] for a proof). For the sake of conciseness, we identify this forest with its edge set F_i. The partition of E into these forests is called the *forest decomposition* of $<$. For vertices $v < w$, we say v is *left* of w. If there is an edge between v and w, we call this a *left-edge* of w.

For any k, we allow to compute k internally vertex-disjoint paths between any two vertices that are contained in a tree T_k of the forest F_k. Hence, throughout the paper, let $s > 1$ be an arbitrary but fixed vertex of G and let k be a positive integer that is at most the number of left-edges of s. The vertex s will be the start vertex of the k internally vertex-disjoint paths to find (the end vertex will be left of s). E.g., if we choose s as the last vertex of the MAO (or any other vertex with at least that many left-edges), k can be chosen as any value that is at most the degree of vertex n; in particular, k can be chosen arbitrary in the range $1, \ldots, \delta$, as claimed in the introduction.

For $i \in \{1, \ldots, k\}$, let T_i the component of F_i that contains s. As $i \leq k$, T_i is a tree on at least two vertices. Let the smallest vertex r_i of T_i with respect to $<$ be the *root* of T_i. For the purpose of this paper, it suffices to consider the subgraph of G induced by the edges of $T_1, \ldots, T_k$.

Lemma 1 ([14, Lemma 2.25]). *Let $i \in \{1, \dots, k\}$. Then $V(T_i)$ consists of the consecutive vertices $r_i, r_i + 1, \dots, w$ in $<$ such that $s \leq w$. Moreover, for each vertex $v \in T_i \setminus \{r_i\}$, the vertex set $\{r_i, r_i + 1, \dots, v\}$ induces a connected subgraph of T_i.*

Hence, for every $i \in \{1, \dots, k\}$, every vertex $v > r_i$ of T_i has exactly one left-edge that is in T_i and thus at least i left-edges that are in G. Let $\mathit{left}_i(v)$ be the end vertex of the left-edge of v in F_i. The root r_i of T_i has left-degree exactly $i - 1$, as if it had more, r_i would have a left-edge in F_i and thus not be the root of T_i and, if it had less, the left-degree of $r_i + 1$ cannot be at least i, as this violates the MA Ordering (this uses that G is simple).

We conclude that $r_1 < r_2 \cdots < r_k$. This, the definition of F_i and Lemma 1 imply the following corollary.

Corollary 2. Let $i < j \leq k$ and let v be a vertex with $r_j < v < s$. Then v is in T_j and T_i, $r_i \leq \mathit{left}_i(v) < \mathit{left}_j(v) < v$ and $r_j \leq \mathit{left}_j(v)$.

For a vertex-subset $S \subseteq V$, let $\overline{S} := V \setminus S$. For convenience, we will denote sets $\overline{\{v\}}$ by $\overline{v}$. For a vertex-subset $S \subseteq V$, a set of paths is *S-disjoint* if no two of them intersect in a vertex that is contained in S. Thus, V-disjointness is the usual vertex-disjointness and a set of paths is *$\overline{v}$-disjoint* if every two of them intersect in either the vertex v or not at all.

We represent paths as lists of vertices. The *length* of a path is the number of edges it contains. For a path A, let $\mathit{end}(A)$ be the last vertex of this list and, if the path has length at least one, let $\mathit{sec}(A)$ be the second to last vertex of this list.

A side note on a proof by Mader. In [8, 7], Mader presents a method to find a *pendant pair* $\{v, w\}$ in a simple graph, which is a pair of vertices that is $\min\{d(v), d(w)\}$-connected, where d is the degree function. He chooses an inclusion maximal clique and deletes vertices of this clique until it is not maximal anymore. When deleting vertices, new edges are added to preserve the degree of vertices that are not in the clique. Next the clique is enlarged to a maximal one; this procedure is iterated until every edge has an incident vertex in the clique. Then a pendent pair of the original graph can be found. It turns out that this method is a preliminary version of MAOs: The order in which the vertices are added to the clique is in fact a maximal adjacency ordering, and, for every k-connected pendent pair found by this method, there is a MAO whose forest decomposition contains this pair as the end vertices of an edge in F_k. This seems to be widely unknown and we are only aware of one place in literature where this similarity is (briefly) mentioned [9, p. 443]. Mader's existential proof can in fact be made algorithmic. However, MAOs provide a much nicer structure, as they work directly on the original graph.

3 The Loose Ends Algorithm

We first consider the slightly weaker problem of computing k internally vertex-disjoint paths between s and the root set $\{r_1, \ldots, r_k\}$. We will extend this to compute k internally vertex-disjoint paths between two vertices in the next section.

Lemma 3. *Algorithm 1 computes k $\overline{s}$-disjoint paths in $T_1 \cup \cdots \cup T_k$ from s to $\{r_1, \ldots, r_k\}$ in time $O(|E(T_1 \cup \cdots \cup T_k)|) \subseteq O(n+m)$.*

The outline of our algorithm is as follows. We initialize each A_i to be the path that consists of the two vertices s and $\mathit{left}_i(s)$ (in that order). The vertices $\mathit{left}_i(s)$ are marked as *active*; throughout the algorithm, let a vertex be *active* if it is an end vertex of an unfinished path A_i.

So far the A_i are $\overline{s}$-disjoint. We aim for augmenting each A_i to r_i. Step by step, for every active vertex v from $s-1$ down to r_1 in $<$, we will modify the A_i to longer paths, similar as in sweep line algorithms from computational geometry. The modification done at an active vertex v is called a *processing step*. From a high-level perspective, the end vertices of several paths A_i may be replaced or augmented by new end vertices w such that $r_i \leq w < v$ during the processing step of v. Such vertices w are again marked as active, which results in a continuous modification of each A_i to a longer path. By the above restriction on w, each path A_i will have strictly decreasing vertices in $<$ throughout the algorithm. At the end of the processing step of v, we unmark v from being active.

Let v be the active vertex that is largest in $<$. Assume that v is the end vertex of exactly one A_i. If $v = r_i$, A_i is finished. Otherwise, we append the vertex $\mathit{left}_i(v)$ to A_i (see Algorithm 1). The important aspect of this approach is that the index of the path A_i predetermines the vertex that augments A_i. Clearly, this way A_i will reach r_i at some point, according to Lemma 1.

However, if at least two paths end at v, this approach does not ensure vertex-disjointness. Let $A_{j_1}, \ldots, A_{j_l}$ these $l \geq 2$ paths and assume $j_1 < j_2 < \cdots < j_l$. We first replace the end vertex v of A_{j_i} with the vertex $\mathit{left}_{j_{i-1}}(\mathit{sec}(A_{j_i}))$ for all $i \neq 1$. We will show that these modified end vertices are strictly smaller than v, which will re-establish the vertex-disjointness. The key idea of the algorithm is then to switch the indices of the l paths appropriately such that the appended vertices are again predetermined by the path index.

Let a *cyclic downshift* on $A_{j_1}, \ldots, A_{j_l}$ replace the index of each path by the next smaller index of a path in this set (where the next smaller index of j_1 is j_l), i.e. we set $A_{j_i} := A_{j_{i+1}}$ for every $i \neq l$ and then replace A_{j_l} with the old path A_{j_1}. We perform a cyclic downshift on $A_{j_1}, \ldots, A_{j_l}$. Note that we did not alter the path A_{j_l} (which was named A_{j_1} before) yet. If $v = r_{j_l}$, A_{j_l} is finished; otherwise, we append the vertex $\mathit{left}_{j_l}(v)$ to A_{j_l}. See Algorithm 1 for a description of the algorithm in pseudo-code. Figure 1 shows a run of Algorithm 1.

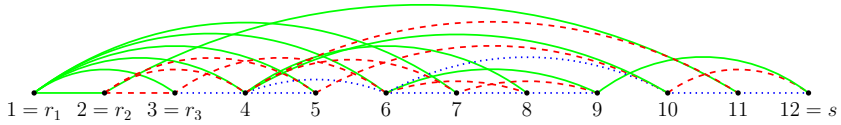

(a) A MAO of a graph G and its forests F_1 (green), F_2 (red, dashed) and F_3 (blue, dotted).

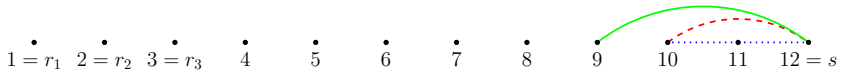

(b) Paths A_1 (green), A_2 (red, dashed) and A_3 (blue, dotted) after the initialization phase and processing vertex 11. The paths A_2 and A_3 end at the largest active vertex 10.

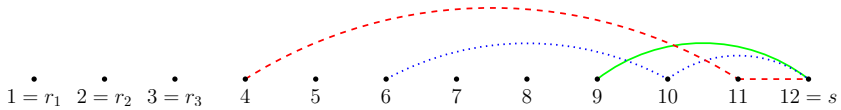

(c) After processing vertex 10, the paths A_2 and A_3 have been shifted, which is here depicted by a color change. The last vertex of A_2 is then replaced, while A_3 is extended in F_3.

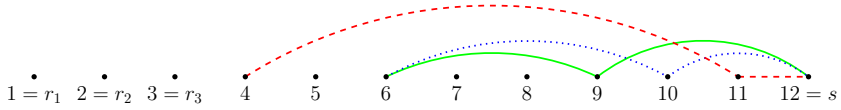

(d) After processing 9, the largest active vertex is 6.

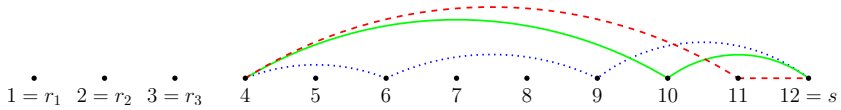

(e) After shifting and extending A_1 and A_3, all three paths meet at the largest active vertex 4.

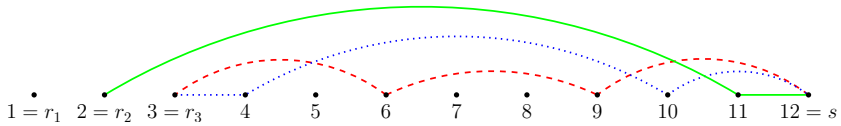

(f) Downshift: The old path A_3 is now A_2, the old A_2 is now A_1 and the old A_1 is now A_3.

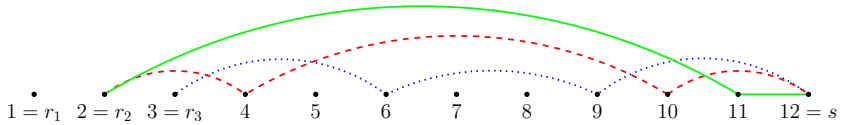

(g) After processing root $r_3 = 3$, A_3 is finished and A_1 and A_2 are shifted.

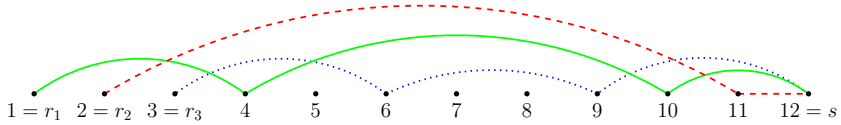

(h) After processing the roots $r_2 = 2$ and $r_1 = 1$, the paths A_2 and A_3 are finished.

Figure 1: A run of Algorithm 1 on the graph depicted in (a) when $s = 12$ and $k = 3$.

Algorithm 1 LooseEnds$(G, <, s, k)$

for all i **do** ▷ initialize all A_i
  $A_i := (s, \mathit{left}_i(s))$
  Mark $\mathit{left}_i(s)$ as active
while there is a largest active vertex v **do** ▷ process v
  Let $j_1 < j_2 < \cdots < j_l$ be the indices of the paths A_{j_i} that end at v
  for $i := 2$ to l **do** ▷ replace end vertices
    Replace $\mathit{end}(A_{j_i})$ with $\mathit{left}_{j_{i-1}}(\mathit{sec}(A_{j_i}))$
    Mark $\mathit{left}_{j_{i-1}}(\mathit{sec}(A_{j_i}))$ as active
  Perform a cyclic downshift on $A_{j_1}, \ldots, A_{j_l}$ ▷ $A_{j_i} := A_{j_{i+1}}$, $A_{j_l} := A_{j_1}$
  if $v = r_{j_l}$ **then**
    A_{j_l} is finished ▷ r_{j_l} is reached
  else
    Append $\mathit{left}_{j_l}(v)$ to A_{j_l} ▷ append predetermined vertex
    Mark $\mathit{left}_{j_l}(v)$ as active
  Unmark v from being active
Output $A_1, \ldots, A_k$

We prove the correctness of Algorithm 1. Before the processing step of any active vertex v, the A_i satisfy several invariants, the most crucial of which are that they are $\{v+1, \ldots, s-1\}$-disjoint and that the vertices of every A_i are decreasing in $<$. In detail, we have the following invariants.

Invariants. *Let $v < s$ be the largest active vertex, or $v := 0$ if there is no active vertex left. Before processing v, the following invariants are satisfied for every $1 \leq i \leq k$:*

(1) The vertices of A_i start with s and are strictly decreasing in $<$.

(2) The path A_i is finished if and only if $\mathrm{end}(A_i) > v$. *In this case,* $\mathrm{end}(A_i) = r_i$.
If A_i is not finished, $r_i \leq \mathrm{end}(A_i) \leq v$ and the last edge of A_i is in T_i.

(3) $\mathrm{sec}(A_i) > v$

(4) Every vertex $w \in A_i$ satisfying $v < w < s$ is not contained in any $A_j \neq A_i$.

(5) $A_i \subseteq T_1 \cup \cdots \cup T_k$

We first clarify the consequences. Invariant (2) implies that the algorithm has finished all paths A_i precisely after processing r_1, and that every A_i ends at r_i. The Invariants (1) and (3) are necessary to prove Invariant (4), which in

turn implies that the A_i are $\{v+1,\dots,s-1\}$-disjoint before processing an active vertex v. Hence, the final paths A_i are $\overline{s}$-disjoint. With Invariant (5) this gives the claim of Lemma 3.

It remains to prove Invariants (1)–(5). Immediately after initializing the paths $A_1,\dots,A_k$, the next active vertex is $\mathit{end}(A_k) < s$. It is easy to see that all five invariants are satisfied for $v = \mathit{end}(A_k)$, i.e. before processing the first active vertex. We will prove that processing any largest active vertex v preserves all five invariants for the active vertex v' that follows v (where $v' := 0$ if v is only remaining active vertex). For this purpose, let A'_i be the path with index i immediately before processing v' and let A_i be the path with index i before processing v; by hypothesis, the A_i satisfy all invariants for v.

For Lines 7 and 13 in the processing step of v, we have to prove the existence of $\mathit{left}_{j_{i-1}}(\mathit{sec}(A_{j_i}))$ and $\mathit{left}_{j_l}(v)$ respectively. In Line 7, we have $i \geq 2$ and $\mathit{end}(A_{j_i}) = v$. Then Invariant (2) implies that A_{j_i} is not finished and $v = \mathit{end}(A_{j_i}) = \mathit{left}_{j_i}(\mathit{sec}(A_{j_i}))$. Thus, $\mathit{left}_{j_{i-1}}(\mathit{sec}(A_{j_i}))$ exists. In Line 13, we have $v \neq r_{j_l}$ and $\mathit{end}(A_{j_l}) = v$ (A_{j_l} before the cyclic downshift). Then Invariant (2) implies that $r_{j_l} \leq v$. This proves $r_{j_l} < v$ and the existence of $\mathit{left}_{j_l}(v)$.

We prove $v' < v$ next. Consider the vertices that are newly marked as active in the processing step of v. According to Line 5 of Algorithm 1, every such vertex is the new end vertex of some path A_{j_i} with end vertex v that was modified in the processing step of v (we do not count index transformations as modifications). There are exactly two cases how A_{j_i} may have been modified, namely either by Line 7 (then $2 \leq i \leq l$ and $\mathit{left}_{j_{i-1}}(\mathit{sec}(A_{j_i}))$ is the vertex that is newly marked as active) or by Line 13 (then $\mathit{left}_{j_l}(v)$ is the vertex that is newly marked as active); in particular, A_{j_i} was not modified by both lines. In the first case, A_{j_i} satisfies Invariant (2) before the processing step of v by hypothesis. In fact, we have $r_{j_i} \leq v$, as $v < r_{j_i}$ implies that A_{j_i} is finished and since $\mathit{end}(A_{j_i}) > v$ would contradict $\mathit{end}(A_{j_i}) = v$.

Hence, the last edge of A_{j_i} is in T_{j_i}, which shows $v = \mathit{left}_{j_i}(\mathit{sec}(A_{j_i}))$. Since $j_{i-1} < j_i$ by Line 5 and due to Corollary 2, we conclude $\mathit{left}_{j_{i-1}}(\mathit{sec}(A_{j_i})) < v$. In the second case, Corollary 2 implies $\mathit{left}_{j_l}(v) < v$. Thus, in both cases, every new active vertex is strictly smaller than v, which proves $v' < v$.

This gives Invariant (1), as every A'_{j_i} starts with s and every new vertex is left of its predecessor in the path by Corollary 2.

For Invariant (2), consider the path A'_i for any i. First, assume that A'_i is finished. Then either A_i is finished or $v = r_i$, according to Line 11 of Algorithm 1 in the processing step of v. In the former case, A_i satisfies Invariant (2) for v and so does A'_i for $v' < v$. In the latter case, we have $v' < v = r_i$ and $\mathit{end}(A'_i) = \mathit{end}(A_{j_1}) = v$.

Second, assume that A'_i was not modified in the processing step of v and is not finished. Then $\mathit{end}(A'_i) < v$, as every path with end vertex at least v is

modified or finished in the processing step of v or finished before. In particular, processing v did not change the index of $A_i = A'_i$. As A_i satisfies Invariant (2) for v by hypothesis, the only condition of Invariant (2) that may be violated for v' is $end(A'_i) < v'$. However, as $end(A'_i) < v$ was marked as active in some previous step of Algorithm 1 and since v' is the largest active vertex, $end(A'_i) \leq v'$. Thus, A'_i satisfies Invariant (2) for v'.

Third, assume that A'_{j_i} was modified in the processing step of v and is not finished. Then A'_{j_i} was modified either by Line 7 or 13. If A'_{j_i} was modified by Line 7, we have $i < l > 1$ after the cyclic downshift, as the path A_{j_1} is not modified by Line 7. In addition, we know $end(A'_{j_i}) = left_{j_i}(sec(A_{j_{i+1}})) < left_{j_{i+1}}(sec(A_{j_{i+1}})) = v$ by Corollary 2 and that the last edge of A'_{j_i} is in T_{j_i}. Thus, $r_{j_i} \leq end(A'_{j_i})$. If A'_{j_i} was modified by Line 13, we have $i = l$ and $r_{j_l} \leq left_{j_l}(v) = end(A'_{j_l})$ by Corollary 2. Then the last edge of A'_{j_l} is in T_{j_l}. In both cases, $end(A'_{j_l})$ is active before processing v' and it follows $end(A'_{j_l}) \leq v'$.

For Invariant (3), assume to the contrary that $sec(A'_i) \leq v'$. Since $v' < v < sec(A_j)$ for all $j \in \{1, \ldots, k\}$, a new end vertex was appended to A'_i in the processing step of v (the end vertex was not replaced, as this would not have changed $sec(A'_i)$). This must have been done in Line 13 of Algorithm 1 and we conclude $v' < v = sec(A'_i)$, which contradicts the assumption.

For Invariant (4), consider Line 7 of the processing step of v. As showed in the proof of $v' < v$ above, we have $left_{j_{i-1}}(sec(A_{j_i})) < v$ for all $1 < i \leq l$. Thus, Invariants (1) and (3) imply that exactly the path A'_{j_l} of the paths $A'_1, \ldots, A'_k$ contains v.

Invariant (5) follows directly from the definition of $left_i$. This concludes the proof of Lemma 3.

Remark. Invariant (4) cannot be strengthened to $A_i \subseteq T_1 \cup \cdots \cup T_i$ or $A_i \subseteq T_i \cup \cdots \cup T_k$, as the cyclic downshifting may force A_i to contain edges of F_k or F_1: During the construction, many active vertices $v \in T_k$ may occur at which all k paths (temporarily) end (see Figure 2, for which each of the three paths A_1, A_2, A_3 contains edges of every tree T_1, T_2, T_3).

So far we have shown an algorithmic proof for the existence of k $\overline{s}$-disjoint paths from s to the roots $r_1, \ldots, r_k$. It remains to show the running time for Lemma 3. Algorithm 1 can be implemented in time $O(n + m)$, as it visits every edge only a constant number of times. Recall that the input graph G does not have to be k-connected; it suffices that G has minimum degree of at least k. A closer look at Algorithm 1 shows that actually only the edges in $T_1 \cup \cdots \cup T_k$ are visited, which gives the slightly better running time $O(|E(T_1 \cup \cdots \cup T_k)|)$.

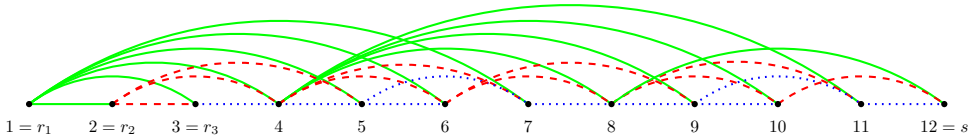

(a) A MAO of a graph H_1 and its forests F_1 (green), F_2 (red, dashed) and F_3 (blue, dotted).

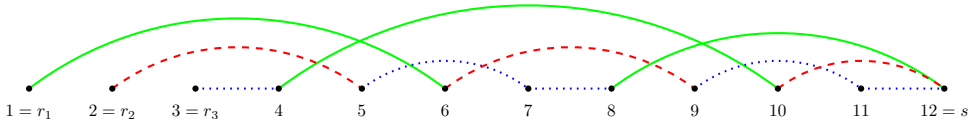

(b) The $\overline{s}$-disjoint paths A_1, A_2 and A_3 in H_1.

Figure 2: The output of Algorithm 1 on the graph depicted in (a) when $s = 12$ and $k = 3$.

4 Computing Vertex-Disjoint Paths Between Two Vertices

We use the algorithm of the last section to prove our following main result.

Theorem 4. *Let $t < s$ be a vertex in T_k. Then k internally vertex-disjoint paths between s and t can be computed in time $O(|E(T_1 \cup \cdots \cup T_k)|) \subseteq O(n + m)$.*

This theorem is directly implied by the following lemma.

Lemma 5. *Let $t < s$ be a vertex in T_k. Then there are k paths $A_1, \ldots, A_k$ with start vertex s and k paths $B_1, \ldots, B_k$ with start vertex t such that $\mathrm{end}(A_i) = \mathrm{end}(B_i)$ for every i and $\{A_1 \cup B_1, \ldots, A_k \cup B_k\}$ is a set of k internally vertex disjoint paths from s to t. Moreover, all paths are contained in $T_1 \cup \cdots \cup T_k$ and can be computed in time $O(|E(T_1 \cup \cdots \cup T_k)|)$.*

A first idea would be to use the loose ends-algorithm twice, once for the start vertex s and once for the start vertex t, in order to find the paths A_i and B_i for all i. However, in general this is bound to fail, as Figure 3 shows. A second attempt may try to finish two paths A_i and B_j whenever they end at the same active vertex. However, this may fail when $i \neq j$, as then two single paths $A_{i'}$ and $B_{j'}$ may remain that end at the respective roots $r_{i'}$ and $r_{j'} > r_{i'}$ such that $B_{j'}$ cannot be extended to $r_{i'}$ without violating the index scheme of Invariant (2), as Figure 4 shows.

We will nevertheless use Algorithm 1 to prove Lemma 5, but in a more subtle way, as outlined next. First, we compute the paths $A_1, \ldots, A_k$ with start vertex s using Algorithm 1, until the largest active vertex v is less or equal t (i.e. the parts of the A_i between s and t are just computed by Algorithm 1). As soon

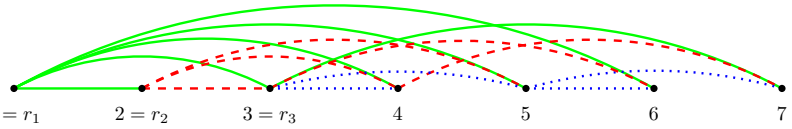

(a) A MAO of a graph H_2 and its forests F_1 (green), F_2 (red, dashed) and F_3 (blue, dotted).

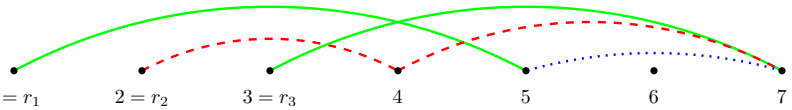

(b) The output of Algorithm 1 on H_2 with start vertex $s = 7$ and $k = 3$ (with a cyclic downshift at 3). Note that each final path A_i ends at r_i due to Invariant (2). The final A_1 contains 5.

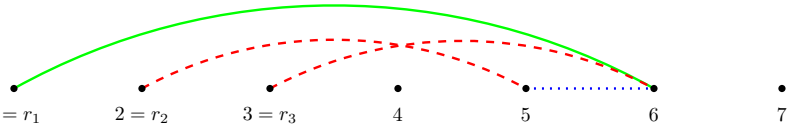

(c) The output of Algorithm 1 on H_2 with start vertex $t = 6$ and $k = 3$. Since B_2 contains 5, both outputs contain the vertex $5 \notin \{s, t, r_1, r_2, r_3\}$, preventing the disjointess of their union.

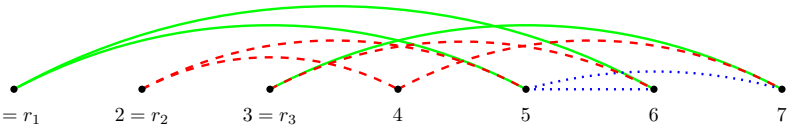

(d) The union of the two outputs. The vertices $s = 7$ and $t = 6$ are separated by $\{3, 5\}$.

Figure 3: A graph H_2 for which applying Algorithm 1 twice does not give internally vertex-disjoint paths.

as $v \leq t$, we additionally construct a second set of paths $B_1, \ldots, B_k$ with start vertex t using Algorithm 1.

The main difference to Algorithm 1 from this point on is that we extend the paths A_i and the paths B_i in parallel (i.e. we take the largest active vertex of both running constructions) such that, after the processing step of v, the vertex v is not contained in any two paths A_i and B_j with $i \neq j$. This ensures the vertex-disjointness.

If no A-path or no B-path ends at v, we again just perform Algorithm 1; then at most one path contains v after the processing step. Otherwise, some A-path and some B-path ends at v. After the processing step at v, we want to have exactly two paths A_j and B_j (i.e. having the same index) that end at v; such a pair of paths is then finished. In order to ensure this, we choose j as the largest index such that A_j *or* B_j ends at v before processing v. If both A_j and B_j end at v, we perform one processing step of Algorithm 1 at v for the A-paths and the

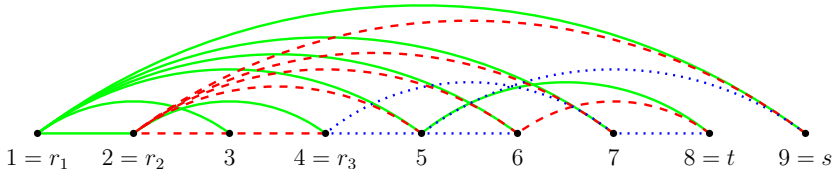

(a) A MAO of a graph H_3 and its forests F_1 (green), F_2 (red, dashed) and F_3 (blue, dotted).

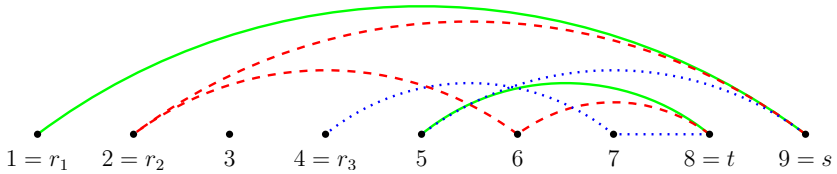

(b) The output of Algorithm 1 on H_3, $k = 3$, with start vertices $s = 9$ and $t = 8$, respectively, such that the path A_3 from s and the path B_1 from t are finished at vertex 5. Note that the path B_3 ends at root $r_3 = 4$ without hitting a matching path from s.

Figure 4: A graph H_3 for which applying Algorithm 1 twice does not give internally vertex-disjoint paths if paths with the same end vertex are finished.

B-paths, respectively, which implies that no other path is ending at v.

Otherwise, exactly one of the paths A_j and B_j ends at v, say A_j. Then B_j is not finished, as we finish only paths having the same index, and the last edge of B_j is in F_j. Hence, there is an index $i < j$ such that B_i ends at v. We then apply a processing step of Algorithm 1 (including a cyclic downshift) on B_j and all B-paths that end at v, and one on all A-paths, respectively. Then the new paths A_j and B_j (due to cyclic downshifts, these correspond to the former A- and B-paths with lowest index ending at v) end at v afterward, but no other A- or B-path, as desired. Note that the replacement of the last edge of (the old) B_j, which did not end at v but, say, at a vertex w, may cause w to be active although neither an A-path nor a B-path ends at w.

For a precise description of the approach, see Algorithm 2. We now show that Algorithm 2 outputs the desired paths and thus proves Lemma 5. The following observations follow directly from Algorithm 2.

Observation 6. *Throughout Algorithm 2 the paths $A_1, \dots, A_k, B_1, \dots, B_k$ satisfy the following properties.*

(1) For every $i \in \{1, \dots, k\}$, A_i and B_i are both finished or both unfinished.

(2) As long as the largest active vertex is larger than t, $B_1 = B_2 = \cdots = B_k = (t)$.

(3) The end vertex of every unfinished path is active.

Algorithm 2 MatchingEnds$(G, <, s, t, k)$ ▷ t is a vertex in T_k, $t < s$

1: **for all** i **do** ▷ initialize all A_i and B_i
2: $A_i := (s, \mathit{left}_i(s))$
3: Mark $\mathit{left}_i(s)$ as active
4: $B_i := (t)$
5: Mark t as active
6: **while** there is a largest active vertex v **do** ▷ process v
7: **if** v=t **then**
8: **for all** i **do** ▷ initialize all A_i
9: **if** $\mathit{end}(A_i) = t$ **then**
10: A_i, B_i are finished
11: **else**
12: Append $\mathit{left}_i(t)$ to B_i
13: Mark $\mathit{left}_i(t)$ as active
14: Unmark t from being active
15: **else**
16: $I_A := \{i | \mathit{end}(A_i) = v\}$
17: $I_B := \{i | \mathit{end}(B_i) = v\}$
18: **if** I_A and I_B are empty **then**
19: Unmark v from being active and go to Line 6
20: $j := \max(I_A \cup I_B)$
21: **for all** pairs (i_1, i_2) of consecutive indices $i_1 < i_2$ in $I_A \cup \{j\}$ **do** ▷ replace ends
22: Replace $\mathit{end}(A_{i_2})$ with $\mathit{left}_{i_1}(\mathit{sec}(A_{i_2}))$
23: Mark $\mathit{left}_{i_1}(\mathit{sec}(A_{i_2}))$ as active
24: **for all** pairs (i_1, i_2) of consecutive indices $i_1 < i_2$ in $I_B \cup \{j\}$ **do** ▷ replace ends
25: Replace $\mathit{end}(B_{i_2})$ with $\mathit{left}_{i_1}(\mathit{sec}(B_{i_2}))$
26: Mark $\mathit{left}_{i_1}(\mathit{sec}(B_{i_2}))$ as active
27: Perform a cyclic downshift on all A_i with $i \in I_A \cup j$
28: Perform a cyclic downshift on all B_i with $i \in I_B \cup j$
29: **if** $v = \mathit{end}(A_j) = \mathit{end}(B_j)$ **then** ▷ if and only if $I_A \neq \emptyset \neq I_B$
30: A_j, B_j are finished
31: **else if** $v = \mathit{end}(A_j)$ **then**
32: Append $\mathit{left}_j(v)$ to A_j ▷ append predetermined vertex
33: Mark $\mathit{left}_j(v)$ as active
34: **else if** $v = \mathit{end}(B_j)$ **then**
35: Append $\mathit{left}_j(v)$ to B_j ▷ append predetermined vertex
36: Mark $\mathit{left}_j(v)$ as active
37: Unmark v from being active
38: Output $A_1, \ldots, A_k$, $B_1, \ldots, B_k$

We prove Lemma 5 by showing that Algorithm 2 outputs the desired paths. Before the processing step of any active vertex v, the paths A_i and B_i satisfy several invariants, the most crucial of which are that they are $\{v+1, \dots, s-1\} \setminus \{t\}$-disjoint and that the vertices of every A_i and B_i are decreasing in $<$. We will prove the following invariants.

Invariants. *Let $v < s$ be the largest active vertex, or $v := 0$ if there is no active vertex left. Before processing v, the following invariants are satisfied for every $1 \leq i \leq k$:*

(1) A_i starts with s, B_i starts with t, and the vertices of both paths are strictly decreasing in $<$.

(2) The paths A_i and B_i are finished if and only if $v < \mathrm{end}(A_i) = \mathrm{end}(B_i)$. If A_i and B_i are not finished, then $r_i \leq \mathrm{end}(A_i) \leq v$, $r_i \leq \mathrm{end}(B_i) \leq v$, and the last edge of A_i as well as the last edge of B_i (if B_i has length at least 1) are in T_i.

(3) $\mathrm{sec}(A_i) > v$. If $v \geq t$, $B_i = (t)$. If $v < t$, either B_i is finished with $B_i = (t)$ or B_i has length at least 1 such that $\mathrm{sec}(B_i) > v$.

(4) Let $w \neq t$ be a vertex with $v < w < s$. If $w \in A_i \cup B_i$, w is neither contained in a path $A_j \neq A_i$ nor in a path $B_j \neq B_i$. If $w \in A_i \cap B_i$, A_i and B_i are finished with $w = \mathrm{end}(A_i) = \mathrm{end}(B_i)$.

(5) $A_i \cup B_i \subseteq T_1 \cup \dots \cup T_k$

Before proving these invariants, we first clarify some of their consequences. Invariant (2) implies that the algorithm has finished all paths when $v = 0$ and that the end vertices of A_i and B_i match for all i. Invariants (1) and (3) will be necessary to prove Invariant (4), which in turn implies that the paths $A_1 \cup B_1, \dots, A_k \cup B_k$ are internally vertex-disjoint. Invariant (5) settles the first part of the second claim of Lemma 5. We continue with further consequences of some of these invariants, which will be used later.

Observation 7. *Let $v < s$ be the largest active vertex, or $v := 0$ if there is no active vertex left. Before processing v, we have the following observations:*

(1) Assume Invariants (1) and (3). Then, for every $1 \leq i \leq k$, all vertices of the paths A_i and B_i except $\mathrm{end}(A_i)$ and $\mathrm{end}(B_i)$ are greater than v before processing v.

(2) Assume Invariant (2). Then no finished path is modified while processing v, as Algorithm 2 modifies A_i or B_i, $1 \leq i \leq k$, only if at least one of them ends at v.

(3) *Assume Invariants (2) and (3). Then the largest active vertex after processing $v > 0$ is smaller than v.*

The following proof uses the fact that new active vertices are only created after Lines 22 and 25. When Line 22 replaces the end vertex of A_{i_2}, the new end vertex $\mathrm{left}_{i_1}(\mathrm{sec}(A_{i_2}))$ *is smaller than the old one due to Corollary 2. By the previous observation, A_{i_2} is not finished. According to Invariant (2), the new end vertex of A_{i_2} is smaller than v.*

When invoking Line 25, $|I_B| \geq 2$ and, hence, there is at least one path B_i that ends at $v \neq t$: According to Invariant (3), $v < t$ and B_{i_2} has length at least one such that $\mathrm{sec}(B_i) > v$*. Thus, the new end vertex of B_{i_2} is smaller than the old one and, in particular, smaller than v.*

We now prove Invariants (1)–(5) for all largest active vertices encountered in the algorithm. Immediately after initializing the paths $A_1, \dots, A_k$ and $B_1, \dots, B_k$, the largest active vertex is $end(A_k) < s$ or t. It is easy to see that all five invariants are satisfied for $v = \max\{end(A_k), t\}$, i.e. before processing the first active vertex. We will prove that processing any largest active vertex v preserves all five invariants for the largest active vertex v' after having processed v (here, we set $v' := 0$ if v is the last active vertex). For this purpose, let A_i and B_i be the paths with index i immediately before processing v and let A'_i and B'_i be the paths with index i immediately before processing v'. By induction hypothesis, A_i and B_i satisfy the desired invariants for the vertex v; hence, we can also use their implications in Observation 7. In particular, $v' < v$ due to Observation 7.(3).

We assume first $v \neq t$ and that v is not the end vertex of any of the paths $A_1, \dots, A_k, B_1, \dots, B_k$ (the following cases can thus assume that at least one such path ends at v). Then Line 19 cancels the processing step without changing any path. According to Observation 7.(1), v is not contained in any of the paths $A_1, \dots, A_k, B_1, \dots, B_k$. The paths $A'_1 = A_1, \dots, A'_k = A_k$, $B'_1 = B_1, \dots, B'_k = B_k$ clearly satisfy Invariants (1) and (5) also for v'. They also satisfy Invariant (3) for v', since $v' \geq t$ if $v > t$, and Invariant (2) for v', since the largest end vertex of any unfinished path is at most v', according to Observation 6.(3). By Invariant (2) for v and Observation 7.(1), any vertex $w \neq t$ with $v' < w \leq v$ is contained in A_i or B_i only if A_i or B_i is unfinished and has end vertex w. Then w is active due to Observation 6.(3), which contradicts the choice of v'. Any $w > v$ would give Invariant (4) by induction hypothesis. Thus, Invariant (4) is satisfied for v'.

Next, we assume $v > t$. Then the paths $A_1, \dots, A_k$ are all not finished and, according to Invariant (3), $B_1 = \dots = B_k = (t)$. We thus have $I_B = \emptyset$ in Line 17 and $j = \max I_A$ in Line 20. In particular, the paths $B_1, \dots, B_k$ are not modified while processing v and the changes of $A_1, \dots, A_k$ are identical to the ones in the loose ends algorithm. Thus, $A'_1, \dots, A'_k$ satisfy the invariants of the loose ends algorithm for v', which imply the invariants of Algorithm 2 for v', as $t \leq v'$.

We assume $v = t$. Then $\mathit{left}_i(t)$ exists in Line 12, since we have $\mathit{end}(A_i) \neq t$, which implies by Invariant (2) that A_i is not finished and $r_i \leq \mathit{end}(A_i) < t = v$. The paths $A'_1, \ldots, A'_k, B'_1, \ldots, B'_k$ satisfy Invariant (1) for v' due to $\mathit{left}_i(t) < t$ for every i, Observation 6.(2) and Invariant (1) for v. We prove Invariant (2) for v': If A'_i and B'_i are finished, $\mathit{end}(A_i) = t$ before processing v and we conclude $v' < \mathit{end}(A'_i) = \mathit{end}(B'_i) = t$. Otherwise, the algorithm leaves both A'_i and B'_i unfinished by setting $A'_i := A_i$ and $B'_i := (t, \mathit{left}_i(t))$. Since both $\mathit{end}(B'_i) = \mathit{left}_i(t)$ and $\mathit{end}(A'_i) = \mathit{end}(A_i)$ are active after processing t, $r_i \leq \mathit{end}(B'_i) \leq v'$ and $r_i \leq \mathit{end}(A'_i) \leq v'$. In particular, both A'_i and B'_i are of length at least one and have their last edge in T_i, which gives Invariant (2) for v'.

For Invariants (3) and (5), note that $A'_i = A_i$ and $B'_i = (t)$ if B'_i is finished and otherwise $A'_i = A_i$ and $B'_i = (t, \mathit{left}_i(t))$. It remains to prove Invariant (4). If A_i contains a vertex w with $v' < w < t$ (again, any larger w would give Invariant (4) by induction hypothesis), w is the active end vertex of A_i by Observations 7.(1) and 6.(3), which contradicts the choice of v'. Since B'_i consists of t and at most one active vertex that is smaller or equal to v', the vertex w of Invariant (4) does not exist. This proves Invariant (4) for v'.

It only remains to assume $v < t$. According to the statement of our first case, we can additionally assume $I_A \cup I_B \neq \emptyset$ (as defined in Lines 16 and 17). Let j be the index chosen in Line 20 of processing step v. Then Observation 6.(1) ensures that both A_j and B_j are unfinished; hence, the downshifts in Lines 27 and 28 are well-defined. The paths $A'_1, \ldots, A'_k, B'_1, \ldots, B'_k$ satisfy Invariant (1) for v' due to Invariant (1) for v, the fact that $\mathit{left}_i(v) < v$ for all i, and Observation 7.(3).

We prove Invariant (2) for v': By induction, this invariant is true for the paths that were finished before processing v. For all paths A'_i and B'_i that are unfinished after processing v, Observation 6.(3) ensures that $\mathit{end}(A'_i) \leq v'$ and $\mathit{end}(B'_i) \leq v'$. The cyclic downshifts in processing step v imply that the last edge of A'_i and the last edge of B'_i are in T_i for every $i \neq j$. Since A'_j and B'_j are the only paths that may change their status from unfinished to finished during the processing step of v, this gives Invariant (2) for all paths except for A'_j and B'_j. These two paths are finished after processing step v if and only if $v' < v = \mathit{end}(A'_j) = \mathit{end}(B'_j)$. In this case, Invariant (2) for v' is satisfied. The case $\mathit{end}(A'_j) \neq v = \mathit{end}(B'_j)$ (and, by symmetry, the case $\mathit{end}(A'_j) = v \neq \mathit{end}(B'_j)$) only occurs if $I_A = \emptyset$ in processing step v. In this case $A'_j = A_j$ and $\mathit{end}(B'_j) = \mathit{left}_j(v)$. By Observation 6.(3) and Invariant (2) for v, $r_j \leq \mathit{end}(A'_j) \leq v'$ and the last edge of A_j is in T_j. This gives Invariant (2) for v'.

In order to prove Invariant (3) for v', observe that all modified paths are unfinished and have length at least one. Thus, we only have to prove that, when appending a vertex to a path in processing step v, the new second last vertex (i.e. the old end vertex) is greater than v'. The algorithm does this only in Lines 32 and 35, in which $v > v'$ is the new second last vertex.

We prove Invariant (4) for v'. If $v < w < s$, this follows directly from Invariant (4) for v, so let $v' < w \leq v$. First, suppose $w < v$ such that w is contained in $A'_i \cup B'_i$. Then w is active after processing v by Invariants (1) and (3) for v', which contradicts the choice of v'. Second, suppose $w = v$. Let j be the index chosen in Line 20 of processing step v. If both paths A'_j and B'_j contain v, both end at v and are finished by Line 30. It remains to prove that $v = w$ is not contained in any other path than A'_j and B'_j. If any path A_i or B_i contains v before processing step v, it contains v as end vertex by Observation 7(1). All paths with v as end vertex, except for the one with smallest index, get a new end vertex in processing step v. After the cyclic downshift in processing step v, only the paths with index j contain v. This proves Invariant (4) for v'.

Invariant (5) follows straight from the definition of left_i. This concludes the proof of Invariants (1)–(5) for every $v < s$.

As in the loose ends algorithm, the running time of Algorithm 2 is upper bounded by $O(|E(T_1 \cup \cdots \cup T_k)|)$ and thus by $O(n + m)$, as it suffices to visit every edge in these trees $T_1, \ldots, T_k$ only a constant number of times.

4.1 Variants

Several variants of Menger's theorem [11] are known. Instead of computing k paths between two vertices, we can compute paths between a vertex and a set of vertices (*fan variant*) and between two sets of vertices (*set variant*). Our algorithm extends to these variants.

Theorem 8. *Let G be a simple graph and $<$, s and $T_1, \ldots, T_k$ be defined as in Section 2.*

(i) (Fan variant) Let $T = \{t_1, \ldots, t_k\}$ be a subset of V such that $r_i \leq t_i < s$ for every i. Then k internally vertex-disjoint paths between s and T can be computed in time $O(|E(T_1 \cup \cdots \cup T_k)|) \subseteq O(n + m)$.

(ii) (Set variant) Let $T = \{t_1, \ldots, t_k\}$ and $S = \{s_1, \ldots, s_k\}$ be disjoint vertex sets such that $r_i \leq t_i < s$ and $r_i \leq s_i \leq s$ for every i. Then k internally vertex-disjoint paths between S and T can be computed in time $O(|E(T_1 \cup \cdots \cup T_k)|) \subseteq O(n + m)$.

Proof. Let G' be the subgraph of G that is induced by the vertex set $\{1, \ldots, s\}$. Clearly, $<$ is also a MAO of G' (restricted to the vertices $1, \ldots, s$) and the relevant parts of the trees $T_1, \ldots, T_k$ are preserved by Lemma 1. For (i), augment G' by a new vertex $s + 1$ with k edges to $t_1, \ldots, t_k$; this preserves that $<$ is a MAO. Applying Theorem 4 to this graph on the vertices s and $s + 1$ gives the claim. For (ii), we augment G' by two new vertices of degree k with neighborhoods S and T, respectively, and apply Theorem 4 on the vertices $s + 1$ and $s + 2$. In both cases, the running time is not increased. □

Mixed Connectivity. Let $G = (V, E)$ be a multigraph and let $\alpha : V \to \mathbb{N}^+$ be a weight function on its vertices. A set of paths connecting two vertices s and t of G is called α-*independent* if every vertex $v \notin \{s, t\}$ is contained in at most $\alpha(v)$ of these paths. Let a multigraph be α-*simple* if the number of edges between every two vertices v and w is at most $\min\{\alpha(v), \alpha(w)\}$.

For α-simple multigraphs G, Nagamochi [12] generalized the existential variant of Theorem 4 by showing that there are k α-independent s-t-paths, where s and t are chosen as in Theorem 4. It is possible to modify Algorithm 2 to compute also these paths without increasing its running time, by replacing the two cyclic downshifts by a more complicated algorithm to transform the path indices.

Acknowledgments. We wish to thank Solomon Lo for pointing out a connection between MAOs and Mader's proof about pendant pairs.

References

[1] S. R. Arikati and K. Mehlhorn. A correctness certificate for the Stoer-Wagner min-cut algorithm. *Information Processing Letters*, 70(5):251–254, 1999.

[2] R. Diestel. *Graph Theory*. Springer, fourth edition, 2010.

[3] S. Even and R. E. Tarjan. Network flow and testing graph connectivity. *SIAM Journal on Computing*, 4(4):507–518, 1975.

[4] A. Frank. On the edge-connectivity algorithm of Nagamochi and Ibaraki. Laboratoire Artemis, IMAG, Université J. Fourier, Grenoble, March 1994.

[5] M. R. Henzinger. A static 2-approximation algorithm for vertex connectivity and incremental approximation algorithms for edge and vertex connectivity. *Journal of Algorithms*, 24:194–220, 1997.

[6] A. V. Karzanov. O nakhozhdenii maksimal'nogo potoka v setyakh spetsial'nogo vida i nekotorykh prilozheniyakh (in Russian; On finding a maximum flow in a network with special structure and some applications). *Matematicheskie Voprosy Upravleniya Proizvodstvom*, 5:81–94, 1973.

[7] W. Mader. Existenz gewisser Konfigurationen in n-gesättigten Graphen und in Graphen genügend großer Kantendichte. *Mathematische Annalen*, 194:295–312, 1971.

[8] W. Mader. Grad und lokaler Zusammenhang in endlichen Graphen. *Mathematische Annalen*, 205:9–11, 1973.

[9] W. Mader. On vertices of degree n in minimally n-connected graphs and digraphs. *Bolyai Society Mathematical Studies (Combinatorics, Paul Erdős is Eighty, Keszthely, 1993)*, 2:423–449, 1996.

[10] R. M. McConnell, K. Mehlhorn, S. Näher, and P. Schweitzer. Certifying algorithms. *Computer Science Review*, 5(2):119–161, 2011.

[11] K. Menger. Zur allgemeinen Kurventheorie. *Fundamenta Mathematicae*, 10:96–115, 1927.

[12] H. Nagamochi. Sparse connectivity certificates via MA orderings in graphs. *Discrete Applied Mathematics*, 154(16):2411–2417, 2006.

[13] H. Nagamochi and T. Ibaraki. Computing edge-connectivity in multigraphs and capacitated graphs. *SIAM Journal on Discrete Mathematics*, 5(1):54–66, 1992.

[14] H. Nagamochi and T. Ibaraki. *Algorithmic Aspects of Graph Connectivity*. Cambridge University Press, 2008.

[15] J. M. Schmidt. Contractions, removals and certifying 3-connectivity in linear time. *SIAM Journal on Computing*, 42(2):494–535, 2013.

[16] M. Stoer and F. Wagner. A simple min-cut algorithm. *Journal of the ACM*, 44(4):585–591, 1997.

[17] R. E. Tarjan and M. Yannakakis. Simple linear-time algorithms to test chordality of graphs, test acyclicity of hypergraphs, and selectively reduce acyclic hypergraphs. *SIAM Journal on Computing*, 13(3):566–579, 1984.

[18] H. Whitney. Non-separable and planar graphs. *Transactions of the American Mathematical Society*, 34(1):339–362, 1932.